Analysis of a Multiple Content Variant Extension of the Multimedia Broadcast/Multicast Service

George Xylomenos, Konstantinos Katsaros and George C. Polyzos
{xgeorge,ntinos,polyzos}@aueb.gr
Mobile Multimedia Laboratory
Department of Informatics
Athens University of Economics and Business
Patision 76, Athens 104 34, Greece

***Abstract*— This paper describes and analyzes an extension of the Multimedia Broadcast / Multicast Service (MBMS) that supports the distribution of multiple variants of the same content to heterogeneous receivers. We first outline the standard MBMS model and then describe our extended MBMS model, detailing the modifications that it imposes on MBMS state management and signaling procedures. We then provide an analytical framework for the comparison of our extended model with both a single and a multiple service approach based on standard MBMS mechanisms. We apply this framework to a practical setting and show that our proposal is far more scalable in terms of signaling overhead than its standards based competitor.**

I. Introduction

Cellular systems are increasingly becoming feasible platforms for multimedia services, due to the high bandwidth of 3rd Generation (3G) systems. While the amount of bandwidth consumed by services such as video distribution makes them too expensive for most users, costs can be dramatically reduced when many users desire to receive the same service, by transmitting the corresponding data only once per cell. This can be achieved either by *broadcast*, where all users receive the service, or by *multicast*, where only selected users receive the service. Similar reasoning in the past led to the introduction of IP multicasting into the Internet. Recently, the *Universal Mobile Telecommunications Systems* (UMTS), specified by the *3rd Generation Partnership Project* (3GPP), introduced the *Multimedia Broadcast/Multicast Service* (MBMS) in its Release 6 specifications. MBMS allows resource sharing throughout the UMTS network, including over the air interface [3]; it is suitable for applications as diverse as media streaming and file downloads [5].

In this paper we concentrate on MBMS multicasting, as it is expected to be more important commercially than MBMS broadcasting. The reason is that multicasting is more suitable for commercial applications targeting a specific set of, possibly paying, subscribers. Both IP and MBMS multicasting deliver the same content to all receivers, so as to transmit data only once over each link. When the receivers are heterogeneous however, for example, terminals with different screen resolutions or users with different budgets, it is difficult to select a variant of the content that will satisfy everyone. If a high quality, expensive, variant is distributed, some terminals will not be able to receive it and some users will not want to pay for it; if a low quality, cheap, variant is distributed, some users will be unsatisfied.

In our research for the IST B-Bone project we have extended the MBMS model so as to support the distribution of multiple variants of the same content to different receivers. The desired variant can be dynamically selected by each receiver, based on terminal capabilities and/or user preferences. Our approach can be combined with layered coding in order to transmit only a single variant over each link between the sender and the receivers.

In this paper, we first describe how our model extends the standard MBMS model, and then analytically evaluate our model against two standards based alternatives. In Section 2 we introduce the standard MBMS model whereas in Section 3 we describe our extended MBMS model, along with its modifications to the standard state management and signaling procedures. In Section 4 we provide an analytical framework for the comparison of our extended model against an alternative approach based on standard MBMS mechanisms and in Section 5 apply this framework in order to show the performance gains offered by our model. Finally, in Section 6 we summarize our conclusions and discuss future work.

II. The standard MBMS model

The functional entities of a UMTS network that are affected by MBMS are shown in Figure 1. The *Broadcast/Multicast Service Center* (BM-SC) is a new entity that controls the services provided by MBMS. The *Gateway GPRS Support Node* (GGSN), the *Serving GPRS Support Node* (SGSN), the *Radio Access Network* (RAN) and the *User Equipment* (UE) are existing UMTS components that must be modified for MBMS. The content sources are beyond the scope of MBMS standardization; they may be internal or external to the network.

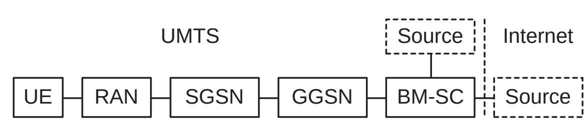

Fig. 1. Components of MBMS.

In the original IP multicasting model, each multicast group is identified by a class D IP address. Any host can *join* the multicast group to start receiving packets sent to it and later *leave* the group to stop receiving such packets. Any host can send packets to the group, even non group members [6]. This *open group* model is clearly not very attractive for content providers. The MBMS multicasting model departs from IP multicasting in that each multicast group is identified by a class D IP address *and* an *Access Point Name* (APN), which effectively identifies a specific GGSN [4]. A more significant difference is that a *closed group* model is used. First, a UE must first subscribe to a group in order to be allowed to later join it, using a separate subscription mechanism. Second, only the GGSN identified by the APN may send to the group; data are first processed by the BM-SC and then delivered by the GGSN. This model enables the provision of commercial services over MBMS [3].

The IP multicasting implementation is split into *local* mechanisms, which track group membership, and *global* mechanisms, which route multicast packets between networks. The only local mechanism defined is the *Internet Group Management Protocol* (IGMP) [6], a query/response protocol suitable for Ethernets; many global mechanisms have been proposed, each using a different routing protocol. In MBMS multicasting, the local mechanisms cover the entire network served by a GGSN: the GGSN acts as the interface between the UMTS network, where MBMS multicasting is used, and the Internet, where IP multicasting is used. The GGSN may however receive content for distribution to a group via either IP multicast or unicast.

While the 3GPP has defined a rich *Quality of Service* (QoS) model for UMTS networks, with multiple service classes and parameters, MBMS services are required to use the same QoS parameters for the entire distribution tree, as defined by the BM-SC. This simplifies tree maintenance, as all paths must support the same QoS, and eliminates the need for QoS negotiations with each UE.

In order to support MBMS multicasting, each node in a UMTS network must maintain two types of state for each multicast group. First, packet forwarding state is required so that the node may determine which of its children should receive a packet; this state is kept on a per group basis. Second, user accounting state is required so that the network may charge the receivers; this state is kept on a per group and a per UE basis. Each node maintains a *MBMS Bearer Context* (MBC) for each multicast group and a *MBMS UE Context* (MUEC) for each UE that is currently a member of the group [4], as shown in Figure 2 (only fields relevant to our work are included).

The MBC contains information for the entire group, such as its IP multicast address and its QoS parameters. The MBC also includes a table indicating which downstream nodes should receive packets addressed to that group. For example, in Figure 2, child #1 should receive packets but child #2 should not. When a multicast packet arrives, the node examines the MBC and forwards the packet to each child marked 1. The MUEC contains information for a UE served by the node that is currently a member of the group. It is created (destroyed)

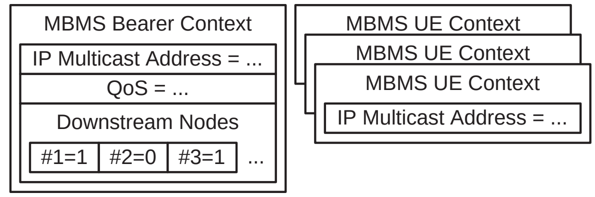

Fig. 2. MBMS contexts in the standard model.

when the UE joins (leaves) the group. Each MUEC is linked to a MBC via its IP multicast address. When the node forwards data to a multicast group, it uses the MUECs linked to the MBC to charge the UEs.

When a MBMS multicast service is to be offered, the data describing it, such as IP multicast address, APN and QoS parameters, are entered into a new MBC at the BM-SC. Additional MBCs and MUECs are dynamically created and destroyed at each node, based on UE initiated signaling. In particular, when a UE desires to join (leave) a group, it sends an IGMP join (leave) message to its GGSN stating the IP multicast address desired [4]. This join/leave mode of IGMP differs from the normal query/response mode used with IP multicasting, but it is actually much more suitable for MBMS [8].

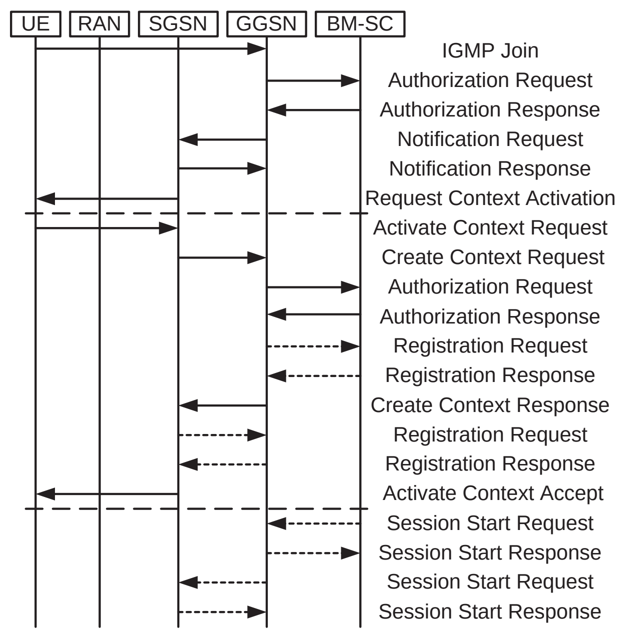

Fig. 3. MBMS Activation, Registration and Session Start procedures.

The IGMP join message triggers the MBMS Activation procedure, shown in Figure 3. The GGSN asks the BM-SC if the UE has subscribed to the group and the BM-SC returns the APN of the GGSN that acts as the source. The GGSN then asks the SGSN if it can handle the MBMS multicasting group. The SGSN responds to the GGSN and notifies the UE that it can proceed with the activation. At this point (first dashed line) the UE knows the APN of the source, which may map to a different GGSN. The UE then requests the SGSN to start sending it multicast data. The SGSN creates the MUEC and notifies the GGSN corresponding to the APN. The GGSN asks the BM-SC if the UE has subscribed to the group, the BM-SC creates the MUEC and responds to the GGSN. The GGSN creates the MUEC and responds to the SGSN, which responds to the UE.

When the first MUEC for a group is created at the GGSN or SGSN, the node initiates the MBMS Registration procedure towards its parent, indicating that it wants to start receiving data addressed to the group. The parent marks the corresponding entry in its MBC with 1 to start forwarding data to that child. Using the information provided in the response, the child creates the MBC for the group. When the service is about to begin transmitting data (second dashed line), the BM-SC informs the GGSN and the GGSN informs each registered SGSN about the properties of the impending transmission with the MBMS Session Start procedure.

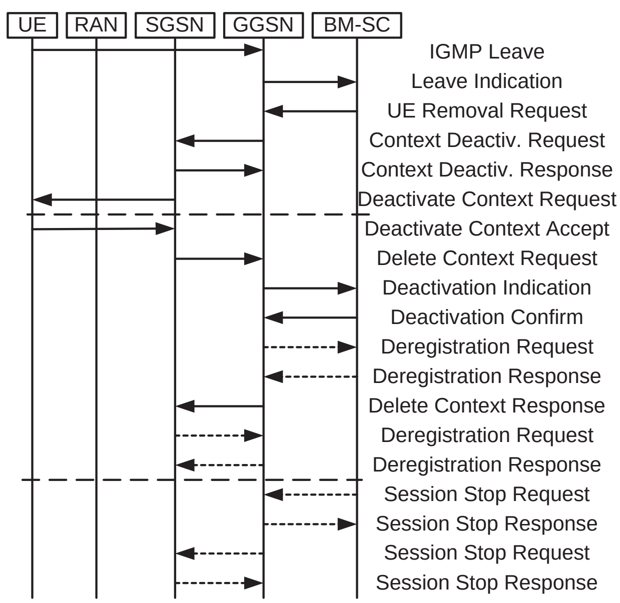

Fig. 4. MBMS Deactivation, Deregistration and Session Stop procedures.

Similarly, an IGMP leave message triggers the MBMS Deactivation procedure, shown in Figure 4. The GGSN informs the BM-SC that the UE is leaving and the BM-SC returns the APN for the group. The GGSN then asks the SGSN to start the deactivation. The SGSN responds to the GGSN and notifies the UE that it can proceed with the deactivation. The UE then requests the SGSN to stop sending it data and the SGSN notifies the GGSN corresponding to the APN. The GGSN destroys the MUEC and notifies the BM-SC. The BM-SC also destroys the MUEC and responds to the GGSN. Finally, the GGSN responds to the SGSN which also destroys the MUEC.

When the last MUEC for a group is destroyed at the GGSN or SGSN, the node initiates the MBMS Deregistration procedure towards its parent, indicating that it wants to stop receiving data for the group and destroys its MBC. The parent marks the corresponding entry in its MBC with 0 to stop forwarding data to that child. When the service finishes transmitting data, the BM-SC informs the GGSN and the GGSN informs each registered SGSN about this with the MBMS Session Stop procedure. Note that a UE can join and leave the group at any time, independently of session start and stop.

III. THE EXTENDED MBMS MODEL

Our extended MBMS model departs from the standard MBMS model by allowing a *single* MBMS service to offer different variants of the same content, providing various trade-offs between bandwidth and quality; we refer to this model as *Multiple Content Variant* (MCV) MBMS. This extension may increase the number of subscribers to a service by satisfying a wider range of heterogeneous receivers, such as terminals with different capabilities or users with different budgets. However, the QoS parameters for all variants of a service remain centrally determined by the BM-SC. The variants must be chosen by the content provider to match common terminals and have sufficiently different costs. For example, for an audio service the variants could be CD, radio and telephone quality sound. The number of available variants must be small to prevent the degeneration of multicast groups to single receivers. We have therefore decided to support up to three variants, numbered 1 (low quality), 2 (medium quality) and 3 (high quality).

Our model requires that the content variants are produced so as to allow a lower quality variant to be derived from a higher quality variant. This enables a node to generate all the variants requested by its children based only on the highest quality variant among them; this is the only variant that the node requests from its parent. This is possible via *layered coding* [7], where the source encodes the lowest quality variant as the base layer and then encodes a series of successive enhancement layers. The next higher quality variant consists of the base layer and the first enhancement layer; each successive variant adds another enhancement layer. The source injects all layers to the multicast distribution tree, and each node forwards to each child only the layers requested by it.

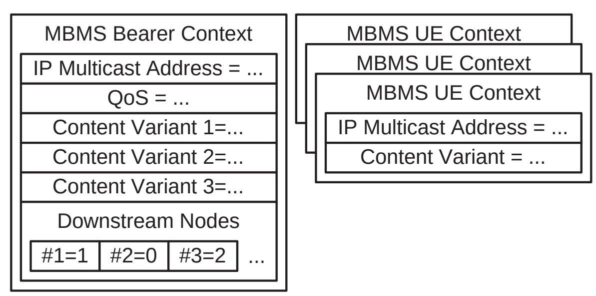

Fig. 5. MBMS contexts in the extended model.

A UE specifies the variant that it wishes to receive in its IGMP join message by including a variant number; it may modify this request at any time by sending a new IGMP join message. For example, the user may request a higher quality audio variant to better hear a passage, or the terminal may request a lower quality audio variant when the bandwidth at its location is limited. Each (internal) node must maintain additional information in the MBC and MUEC, as shown in Figure 5. The MUEC must be extended with the number of the requested variant, allowing the node to charge the UE accordingly. In addition, the downstream nodes table in the MBC must be extended with the number of the variant to forward to each child. For example, in Figure 5, child #1 should receive variant 1, child #2 should not receive anything and child #3 should receive variant 2.

In our extended MBMS model, each node must inform its parent about the variant that it needs to receive, thus allowing the parent to maintain its MBC. The node must thus determine the highest quality variant requested by *any* of its own children; from this variant it may produce any lower quality variants required. We can determine this information by counting the number of MUECs for each content variant of a group. These counters are stored in the MBC, as shown in Figure 5, and each node requests from its parent the highest quality variant with a nonzero counter.

When a UE sends an IGMP join (leave) message, triggering the activation (deactivation) procedure, at the point where a node would create (destroy) a MUEC in the standard MBMS model, in our model the node must instead do one of the following:

- If a MUEC was created, the counter for the corresponding variant is incremented by one.
- If a MUEC was destroyed, the counter for the corresponding variant is decremented by one.
- If a MUEC was modified, the counter for its previous variant is decremented by one and the counter for its current variant is incremented by one.

Another related modification is the addition of a variant number in the request messages of the activation procedure (see Figure 3). Furthermore, at the point where a node would create (destroy) a MBC in the standard MBMS model, in our model the node must instead do one of the following:

- If the first MUEC for a group was created, the MBC is created (one nonzero counter) and the parent is informed to start forwarding the corresponding variant.
- If the counter for a higher quality variant than the one currently requested became nonzero, the parent is informed to send the corresponding variant.
- If the counter for the currently requested variant became zero, the next nonzero counter is located and the parent is informed to send the corresponding variant.
- If the last MUEC for a group was destroyed (all counters are zero), the MBC is destroyed and the parent is informed to stop forwarding all variants.

Another related modification is the addition of a variant number in the request messages of the registration procedure (see Figure 3). Note that new registration messages can be received from a child that is already receiving a group, if the child wishes to change its content variant as UEs join and leave the service or change their content variants. In order to avoid the need for updated session start messages, these messages always include information for *all* available content variants; nodes simply store this information in case it is needed.

IV. Analytical Framework

With the standard MBMS mechanisms, only a single content variant may be supported per service; we refer to this option as Single MBMS. We can however approximate the behavior of MCV MBMS by combining layered coding with a separate MBMS group per layer and having each UE join the groups corresponding to the layers that it needs; we refer to this option as Multiple MBMS. An advantage of this approach is that data packets do not need to indicate the layer that they belong to so as to allow nodes to selectively forward them, as each layer is mapped to a different group. On the other hand, in this approach UEs must join and leave multiple MBMS groups, they must be able to receive multiple MBMS groups simultaneously and there is no guarantee that packets from different layers will be received in a synchronized manner so as to be decoded on time.

Variable	**Description**
N_u	Number of users participating in a service
p_i	Probability that a user requests variant i
S_n	Number of SGSNs in the network
A_m / A_b	Activation messages/bytes (per SGSN)
D_m / D_b	Deactivation messages/bytes (per SGSN)
R_m / R_b	Registration messages/bytes (per SGSN)
DR_m/DR_b	Deregistration messages/bytes (per SGSN)
S_m / S_b	Session Start messages/bytes (per SGSN)
SS_m/SS_b	Session Stop messages/bytes (per SGSN)

TABLE I
List of variables

We will now provide a quantitative comparison of MCV MBMS against Multiple MBMS in terms of the total number and size of the signaling messages exchanged when providing a service with three content variants numbered 1 (low quality), 2 (medium quality) and 3 (high quality). Layered coding is used to generate three content flows which are combined at the receiver to reconstruct the required variant: low quality receivers only need flow 1, medium quality receivers need flows 1 and 2, and high quality receivers need flows 1, 2 and 3. In Multiple MBMS this is performed explicitly, by joining multiple groups. In MCV MBMS this is performed implicitly, by indicating the variant required when joining a group. We use Single MBMS as a baseline; as this option only provides a single content variant, comparing MCV MBMS or Multiple MBMS against it shows the signaling overhead for providing multiple content variants in the extended or in the standard MBMS model.

We omit from our analysis the user plane overhead for two reasons. First, the number of user plane messages is exactly the same in both cases. Second, since in our extended model each user plane message only needs two extra bits, so as to encode the content variant that it belongs to, it is possible to include them in existing fields in order to avoid introducing additional overhead.

We assume that N_u users participate in a multicast MBMS service in a network with a single GGSN and S_n SGSNs.

Procedure	Single MBMS	MCV MBMS	Multiple MBMS
Activation	$A_m N_u$	$A_m N_u$	$A_m(1 + p_2 + 2p_3)N_u$
Deactivation	$D_m N_u$	$D_m N_u$	$D_m(1 + p_2 + 2p_3)N_u$
Registration	$R_m S_n$	$2R_m S_n$	$3R_m S_n$
Deregistration	$DR_m S_n$	$R_m S_n + DR_m S_n$	$3DR_m S_n$
Session Start	$S_m S_n$	$S_m S_n$	$3S_m S_n$
Session Stop	$SS_m S_n$	$SS_m S_n$	$3SS_m S_n$

TABLE II

NUMBER OF SIGNALING MESSAGES REQUIRED.

Procedure	Single MBMS	MCV MBMS	Multiple MBMS
Activation	$A_b N_u$	$A_b N_u + A_m N_u/2$	$A_b(1 + p_2 + 2p_3)N_u$
Deactivation	$D_b N_u$	$D_b N_u$	$D_b(1 + p_2 + 2p_3)N_u$
Registration	$R_b S_n$	$2R_b S_n + R_m S_n$	$3R_b S_n$
Deregistration	$DR_b S_n$	$R_b S_n + R_m S_n/2 + DR_b S_n$	$3DR_b S_n$
Session Start	$S_b S_n$	$S_b S_n + 16 S_m S_n$	$3S_b S_n$
Session Stop	$SS_b S_n$	$SS_b S_n$	$3SS_b S_n$

TABLE III

NUMBER OF SIGNALING BYTES REQUIRED.

Users request each variant with probabilities p_1, p_2 and p_3, where $p_1 + p_2 + p_3 = 1$. All variables used in our analysis are summarized in Table I. To simplify this analysis, we also make the following assumptions:

- All extended MBMS messages indicating a content variant are inflated by 1 byte.
- Users do not change their content variant preferences over time.
- No handovers occur during the period under study.
- At least one receiver of the high quality variant is present in each cell.
- Three assumptions may be made for the number of (de)registrations in our extended model:
  - Best case: The first user to join the service in each cell requests the high quality variant; for deregistrations the last user to leave the service is a high quality user.
  - Worst case: First the low quality, then the medium quality and then the high quality users join each service; for deregistrations this sequence is reversed.

  Since the difference between the best and worst cases turns out to be small, in the comparison tables and graphs given below we show the arithmetic average of the best and worst cases for convenience.

We will separately examine the cost of each MBMS procedure in terms of messages and bytes exchanged. We first point out that with MCV MBMS only half of the exchanged messages (the requests) in the MBMS Activation and Registration procedures need an extra byte to indicate the required content variant; the responses, as well as all messages in the MBMS Deactivation and Deregistration procedures, are as in standard MBMS.

Following the notation presented in Table I, in Single MBMS an MBMS Activation procedure requires $A_m N_u$ messages and $A_b N_u$ bytes to handle all users. In Multiple MBMS this procedure requires $A_m(1 + p_2 + 2p_3)N_u$ messages and $A_b(1 + p_2 + 2p_3)N_u$ bytes, since each user may need to join multiple groups. In MCV MBMS we only need $A_m N_u$ messages, as each user only joins a single group, totaling $A_b N_u + A_m N_u/2$ bytes, due to the extra bytes in the request messages. Similarly, in Single MBMS an MBMS Deactivation procedure requires $D_m N_u$ messages and $D_b N_u$ bytes; the same holds for MCV MBMS, since no extra bytes are needed in this case. In Multiple MBMS this procedure requires $D_m(1 + p_2 + 2p_3)N_u$ messages and $D_b(1 + p_2 + 2p_3)N_u$ bytes, similar to the MBMS Activation case.

The number of messages and bytes required for the MBMS Registration procedure may depend on the number of SGSNs with participants to the service, the content variant preferences of the participants in each cell and the exact ordering of the activations. Since we assumed that at least one high quality receiver exists in each cell, all SGSNs will eventually register with the GGSN for all variants. In Single MBMS, only a single MBMS Registration is performed by each SGSN, therefore $R_m S_n$ messages totaling $R_b S_n$ bytes are needed. In Multiple MBMS on the other hand this procedure requires $3R_m S_n$ messages totaling $3R_b S_n$ bytes. In MCV MBMS, the order of activations is also significant. Under the best case assumption the high quality users will activate first, therefore $R_m S_n$ messages totaling $R_b S_n + R_m S_n/2$ bytes will be required, due to the extra field in the requests; under the worst case assumption the low quality users will join first, then the medium quality ones and then the high quality ones, therefore $3R_m S_n$ messages totaling $3R_b S_n + 3R_m S_n/2$ bytes will be required. The arithmetic average of the two cases is thus $2R_m S_n$ messages totaling $2R_b S_n + R_m S_n$ bytes.

Similarly, in Single MBMS an MBMS Deregistration procedure requires $DR_m S_n$ messages totaling $DR_b S_n$ bytes, while in Multiple MBMS this procedure requires $3DR_m S_n$ messages totaling $3DR_b S_n$ bytes. In MCV MBMS, under the best case assumption the high quality users will deactivate last, therefore $DR_m S_n$ messages totaling $DR_b S_n$ bytes

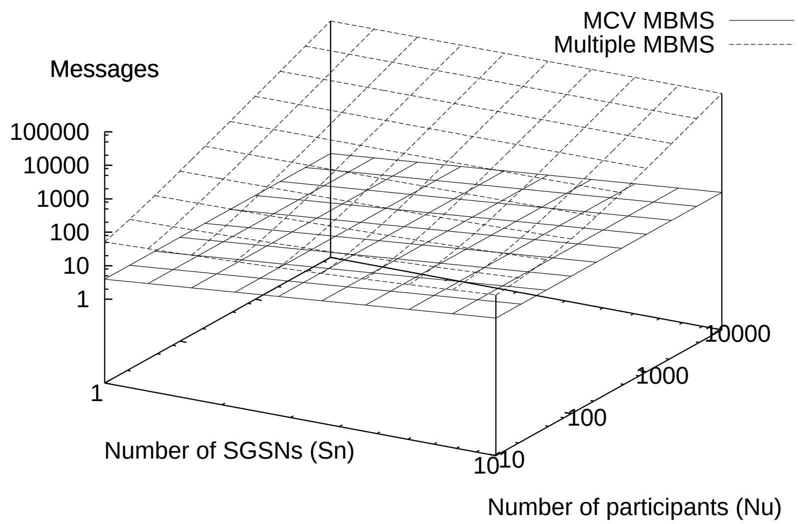

Fig. 6. Additional signaling messages (over Single MBMS) as a function of S_n and N_u.

will be required; under the worst case assumption, the high quality users will deactivate first, then the medium quality ones and then the low quality ones, therefore we will have two registration updates followed by a deregistration, hence $2R_mS_n + DR_mS_n$ messages totaling $2R_bS_n + 2R_mS_n/2 + DR_bS_n$ bytes will be required. The arithmetic average of the two cases is thus $R_mS_n + DR_mS_n$ messages totaling $R_bS_n + R_mS_n/2 + DR_bS_n$ bytes.

The MBMS Session Start procedure also depends on the number of SGSNs with participants to the service, but, due to our assumptions, all SGSNs will have to receive Session Start messages. Therefore, in Single MBMS an MBMS Session Start requires S_mS_n messages totaling S_bS_n bytes; in Multiple MBMS this procedure requires $3S_mS_n$ messages totaling $3S_bS_n$ bytes. In MCV MBMS, since each request message must include the QoS profiles of all available variants, S_mS_n messages totaling $S_bS_n + 16S_mS_n$ bytes are required, assuming that each QoS profile requires 16 bytes and that we need two additional QoS profiles (in addition to the one already included) in the request messages.

Finally, the MBMS Session Stop procedure in Single MBMS requires SS_mS_n messages totaling SS_bS_n bytes; in Multiple MBMS it requires $3SS_mS_n$ messages totaling $3SS_bS_n$ bytes. In MCV MBMS SS_mS_n messages totaling SS_bS_n bytes will be required, since no additional fields are needed. The total number of messages and bytes required per procedure for each option are summarized in Table II and III, respectively.

V. Framework Application

The framework presented above uses variables for the number and sizes of messages required by each signaling procedure, thus allowing the resulting expressions to be used with any set of protocols desired, by simply substituting the message counts and sizes of the protocols under study. We will apply the framework with respect to the *Non Access Stratum* (NAS) protocols, the operation of which is summarized in Figure 3 and 4. The messages shown in the figures between GGSN and SGSN are part of GTP, while the messages between SGSN and UE are part of MBMS-SM. We will also include the IGMP messages between UE and GGSN; while IGMP is not part of MBMS, it is used to trigger Activation and Deactivation and it is executed inside the UMTS network. We omit the Diameter messages between the GGSN and BM-SC, as these are not transported over UMTS protocols.

In order to calculate the number of messages and bytes required, we used the relevant standards for GTP [1], MBMS-SM [2] and IGMP [6] and assumed no errors occurred. We omit the cost of the encapsulating protocols, except for IGMP where we include the cost of the encapsulating IP and GTP packets. We found that $A_m = 8$ and $A_b = 292$, $D_m = 7$ and $D_b = 183$, $R_m = 2$ and $R_b = 59$, $DR_m = 2$ and $DR_b = 50$, $S_m = 2$ and $S_b = 124$, and, finally, $SS_m = 2$ and $SS_b = 50$.

By substituting these values in Table II and adding all rows we find that the total cost in messages for Single MBMS is $15N_u + 8S_n$, for MCV MBMS it is $15N_u + 12S_n$ and for Multiple MBMS it is $15(1+p_2+2p_3)N_u + 24S_n$. Compared to Single MBMS, MCV MBMS introduces extra costs only per SGSN, while Multiple MBMS introduces extra costs per SGSN and per UE. Similarly, by substituting these values in Table III and adding all rows we find that the total cost in bytes for Single MBMS is $475N_u + 283S_n$, for MCV MBMS it is $479N_u + 436S_n$ and for Multiple MBMS it is $475(1 + p_2 + 2p_3)N_u + 849S_n$. Compared to Single MBMS, Multiple MBMS always introduces more overhead than MCV MBMS per SGSN; per UE, Multiple MBMS introduces more overhead than MCV MBMS if $475(1 + p_2 + 2p_3)N_u > 479N_u$, or, if $p_2 + 2p_3 > 0.0084$.

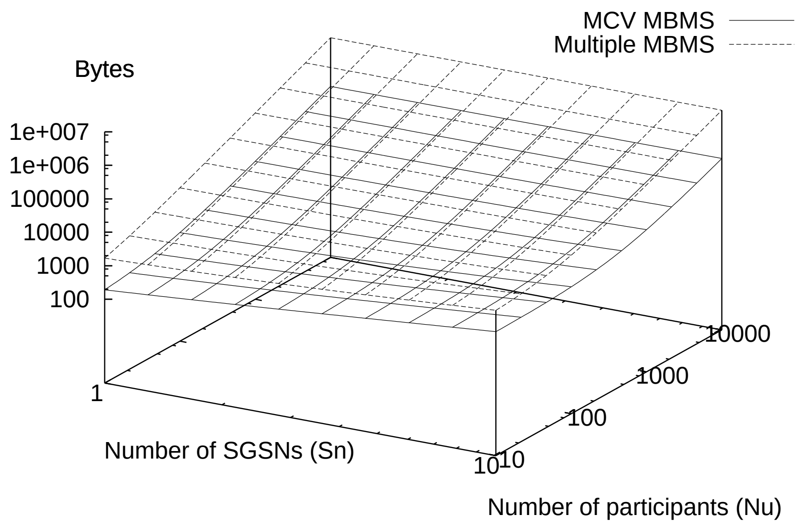

Fig. 7. Additional signaling bytes (over Single MBMS) as a function of S_n and N_u.

To plot these formulas, we assume that the 80-20 rule holds for the content variant preferences: 80% of the users request the low quality variant and the remaining 20% request the other ones; 80% of these users request the medium quality variant and the remaining 20% request the high quality one. Therefore, $p_2 = 0.16$ and $p_3 = 0.04$, and $p_2 + 2p_3 = 0.24 > 0.0084$. With these values, the additional overhead of MCV MBMS over Single MBMS in messages becomes $4S_n$, while for Multiple MBMS it is $3.6N_u + 16S_n$; these are plotted in Figure 6. Similarly, the additional overhead of MCV MBMS over Single MBMS in bytes becomes $4N_u + 153S_n$, while for Multiple MBMS it is $114N_u + 566S_n$; these are plotted in Figure 7. Both figures show results for 10–10,000 participants and 1–10 SGSNs, with all axes logarithmic. In all cases, the overhead of MCV MBMS is far lower than that of Multiple MBMS.

VI. Conclusion and Future Work

In this paper we presented an extended MBMS model supporting the distribution of multiple variants of the same content to heterogeneous receivers. We formulated an analytical framework for the evaluation of the signaling overhead introduced by our MCV MBMS model compared to both Single and Multiple MBMS service options. We then applied this framework to the NAS protocols, showing that MCV MBMS always requires fewer messages than the Multiple MBMS alternative, and also requires fewer bytes even if only 1% of the participants request the medium or high quality variants.

We are currently implementing our MCV MBMS model in the MBMS System Level Simulator developed as part of the IST B-Bone project. This simulator will allow us to study the signaling overhead of each of the alternatives discussed above under dynamic conditions, such as user mobility, random joining and leaving times and random content variant selection, thus relaxing the simplifying assumptions made here.

VII. Acknowledgement

This work was supported by the IST B-Bone project under contract IST-2003-507607.

References

[1] 3rd Generation Partnership Project (3GPP). General Packet Radio Service (GPRS); GPRS Tunneling Protocol (GTP) across the Gn and Gp interface. TS 29.060, V6.6.0, September 2004.

[2] 3rd Generation Partnership Project (3GPP). Mobile radio interface Layer 3 specification; Core network protocols; Stage 3. TS 24.008, V6.6.0, September 2004.

[3] 3rd Generation Partnership Project (3GPP). Multimedia Broadcast / Multicast Service (MBMS); Stage 1. TS 22.146, V6.6.0, September 2004.

[4] 3rd Generation Partnership Project (3GPP). Multimedia Broadcast / Multicast Service (MBMS) user services; Architecture and functional description. TS 23.246, V6.4.0, September 2004.

[5] 3rd Generation Partnership Project (3GPP). Multimedia Broadcast / Multicast Service (MBMS) user services; Stage 1. TS 22.246, V6.2.0, September 2004.

[6] W. Fenner. Internet group management protocol, version 2. RFC 2236, November 1997.

[7] B. Li and J. Liu. Multirate video multicast over the internet: An overview. *IEEE Network*, 17(1):24–29, January/February 2003.

[8] G. Xylomenos. Group management for the Multimedia Broadcast/Multicast Service. In *Proc. of the IST Mobile Summit*, June 2005.